MORAL HAZARD, UNCERTAIN TECHNOLOGIES AND LINEAR CONTRACTS

Martin Dumav, Urmee Khan

Department of Economics
Universidad Carlos III de Madrid, UC Riverside

23 December, 2017

Introduction

Contracting in Uncertain Environments: Simplicity and Robustness

"Real world incentive schemes appear to take less extreme forms than the finely tuned rules predicted by [economic] theory."

Holmstrom and Milgrom (1987)

"Economic guidance can be of help to policymakers in principle. However, the guidance delivered in the form of complex formulae and detailed mathematical characterizations of optimal policies may not be fully appreciated or warmly embraced. ... compelling metrics of "simplicity" are difficult to formulate."

Chu and Sappington (2007)

Contracting with Precise Information: Moral Hazard with Bayesian Actors

'Moral hazard' in a Principal-Agent setting: problems involving unobserved action taken by the agent after the contract is agreed to, that affects the outcome and payoffs

Grossman-Hart (1983) approach: the 'riskiness' of the environment is captured in a common prior probability, so the principal is able to do the following:

-calculate ex-ante expected costs of inducing agent to undertake any available action

-figure out which action she would like the agent to take

-write a contract to 'implement' that action

Contracting with Precise Information: Moral Hazard with Bayesian Actors

The key characteristics of the G-H approach are the following:

- principal knows enough to decide, from an-ex-ante perspective, what action(s) she would like to be implemented

- the information structure allows her to draw inference after the fact, so that implementation via 'reward-punishment' based on some observed signal is feasible

- depending on the details of the model, one can end up with non-linear, even non-monotone contracts (in contract theory terminology, 'high-powered' incentives)

- more often than not, the complicated form the contracts take is not robust to purturbation of the Bayesian information structure

Research Question

- What if the environment itself is not well-understood for it to be modeled as a unique and common prior distribution? (For eg, new technology, new business model, lack of past data, information coming from conflicting sources)

- To what extent can we explain linear contracts as ‘robust’ solutions to moral hazard problems in an environment with such imprecise information?

- What happens to ‘high-powered incentives’ as the information available to the principal becomes more imprecise in comparison to standard Bayesian models?

Some Answers

Main findings:

(1) Building on the approach in Carroll (2015),

- we provide a tractable formulation of moral hazard models with ambiguity

- show that linear contracts are optimal when contracting parties are risk-neutral and use possible worst-case scenarios as guidance for action

(2) Insights about contracts in the presence of ambiguity:

- optimality of linear contracts comes from their ability to align the worst case scenarios for the principal and the agent;

- with imprecise information, ex-post inference becomes less influential, instead ex-ante alignment of perceptions more beneficial

Enviornment : Outcomes and Actions

$\mathcal{Y} \subseteq \mathbb{R}^+$ a compact convex set of outcomes

$\Delta(\mathcal{Y})$ the set of Borel distributions on $\mathcal{Y}$ with the weak* topology

$\mathbb{K}_{\Delta(\mathcal{Y})}$ the class of non-empty, compact, convex subsets of $\Delta(\mathcal{Y})$

The agent chooses an unobservable action a from a compact set $\mathcal{A}$

Each action $a \in \mathcal{A}$ costs $g(a)$ to the Agent

Environment: Information

Principal's information about technology is characterized as a mapping $Q^P(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$

Principal believes that this action induces a set of distributions $Q^P(a) \in \mathbb{K}_{\Delta(\mathcal{Y})}$

Agent's information is characterized by $Q^A(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$

In what follows, we use $Q^i(\cdot)$ to denote both the mapping and its range, a bit of abuse of notation

Environment: Contract and Preferences

A contract is a (continuous) function $w : \mathcal{Y} \to \mathbb{R}^+$ that specifies output contingent payments to the Agent and protects her with limited liability

Each party is risk neutral and evaluates the contract according to its worst possible scenario

Timing

(i) Principal offers a contract w;
(ii) Agent, knowing Q^A, chooses action $a \in \mathcal{A}$;
(iii) output y is realized;
(iv) payoffs are received: $y - w(y)$ to Principal and $w(y) - g(a)$ to Agent.

Benchmark Case: Symmetric Ambiguity

In this version of the model, Principal and Agent has symmetric ambiguous beliefs about the uncertainty: $Q(a) := Q^A(a) = Q^P(a)$ for all $a \in \mathcal{A}$

Contracting Problem - Symmetric Ambiguity

Given symmetric ambiguity Principal uniquely infers: Agent's decision rule

$$a^*(w|Q) = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \left(\min_{q \in Q(a)} \mathrm{E}_q \left[w(y) \right] - g(a) \right)$$

and the best of the worst-case scenarios perceived by the Agent $q_A^*(w|Q) \in Q(a^*(w|Q))$

Contracting Problem with Symmetric Ambiguity

From an arbitrary contract w, Agent's MaxMin value is

$$V_A(w|Q) = \mathsf{E}_{q_A^*(w|Q)}[w(y)] - g(a^*(w|Q))$$

The Principal's guarantee is

$$V_p(w|Q) = \min_{q \in Q(a^*(w|Q))} \mathsf{E}_q[y - w(y)]$$

His worst-case scenarios is $q_P^*(w|Q) \in Q(a^*(w|Q))$

Observation: the Principal and the Agent does not have to agree on the worst-case scenario, i.e., $q_P^*(w|Q)$ and $q_A^*(w|Q)$ can differ

The Principal offers a contract w that maximizes

$$V_p(w|Q) \text{ subject to } V_A(w|Q) \geq 0$$

We assume that contracting relation is viable: for some contract w, we have that $V_p(w|Q) \geq 0$ and $V_A(w|Q) \geq 0$

Optimality of Linear Contracts - Symmetric Ambiguity

A linear contract ℓ is of the form $\ell(y) = \alpha y$ with $\alpha \in (0, 1]$

Theorem (1)

Suppose $Q^A(a) = Q^P(a)$ for all $a \in \mathcal{A}$. There exists a linear contract that maximizes V_p.

Proof Strategy:

Step 1: For any arbitrary (possibly non-linear, non-monotone) viable contract w, can we find a linear contract that does at least as well as w for both the Principal and the Agent?

Yes, we find such a linear contract

The key observation is that linearity ensures that Principal and Agent agrees on the worst case and this helps us to find a tight linear relationship between the P's and A's guarantees.

Step 2: Focusing on the set of linear contracts, one can find an optimal one

Graphical Intuition

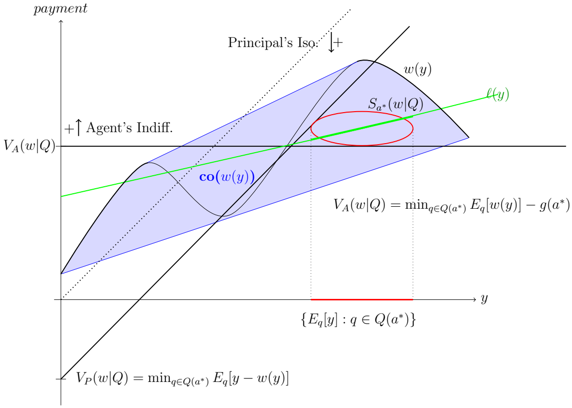

Disagreement on the Worst case

Fix an arbitrary contract w

For any a the Agent's worst case scenario minimizes expected payments received $q_A(w) = \arg\min_{q \in Q(a)} E_q[w(y)]$ and equivalently

$$E_{q_A(w)}[w(y)] \leq E_q[w(y)] \text{ for all } q \in Q(a) \quad (*)$$

In particular, this holds for $q_P(w) \in Q(a)$ – the Principal's worst case that minimizes his expected profits $q_P(w) = \arg\min_{q \in Q(a)} E_q[y - w(y)]$

If the worst cases differ, then the expected payments made are larger than the expected payments received: If $q_A(w) \neq q_P(w)$ then $(*)$ implies $E_{q_A(w)}[w(y)] < E_{q_P(w)}[w(y)]$

Disagreement on the Worst case

Divergence of worst-cases is never beneficial for Principal

For a given contract w, let a be an action that can be chosen by Agent

Agent's worst-case expected payoff lies on the horizontal hyperplane that supports the convex set $\mathcal{H}(a) = \{(E_q[y], E_q[w(y)] : q \in Q(a)\}$ from below (corresponding to Agent's worst-case payoff while choosing a)

For the same action, Principal's worst-case scenario lies on the highest 45^0-degree hyperplane supporting $\mathcal{H}(a)$, representing Principal's iso-profit curve that minimizes expected profit $y - w(y)$ within $\mathcal{H}(a)$

Disagreement on the Worst case

For an arbitrary contract, $\mathcal{H}(a)$ has possibly non-empty interior and hence Principal and Agent do not agree on the worst-case distribution.

$\mathcal{H}(a)$ is contained in an affine convex cone generated by $\{(u, v) \in \mathbb{R}^2_+ | v = \beta\}$ and

$\{(u, v) \in \mathbb{R}^2_+ | u - v = \gamma\}$, where β and γ are Agent's and Principal's worst case payoffs when a is chosen

A linear contract can be found, that reduces $\mathcal{H}(a)$ it to an increasing (one-dimensional) line and aligns the pessimistic expectations

The linear contract yields expected values in the interior of $\mathcal{H}(a)$ and hence improves on w for both parties if the Agent's action choice does not change

If the Agent chooses a different action it will improve her value, and also the Principal's value, since under the linear contract the parties share the common worst-case

Therefore, for any arbitrary contract, for any chosen action we can construct a linear contract that does at least as well.

Linearity and Agreement on the Worst Case

If $w(y)$ is linear then profits $y - w(y)$ is a linear (e.g., monotone increasing) transformation of wages $w(y)$, and hence by the linearity of expected values in probabilities

$$q_P(w) := \arg \min_{q \in Q(a)} E_q[y - w(y)] = \arg \min_{q \in Q(a)} E_q[w(y)] =: q_A(w)$$

Characterzing the Optimal Linear Contract

With a linear contract, $w(y) = \alpha y$, $\alpha \in [0, 1]$the optimal share is found by maximizing

$$\max_a [(1-\alpha) E_{q^*(a)}[y] - \frac{1-\alpha}{\alpha} g(a)] \tag{1}$$

$q^*(a)$, is the common worst-case distribution that minimizes the expected output in $Q(a)$.

For any given interior action, the optimal α that maximizes $(1-\alpha)E_{q^*(a)}[y] - \frac{1-\alpha}{\alpha}g(a)$ is equal to $\alpha = \sqrt{g(a)/E_{q^*(a)}[y]}$.

Using this we can solve for the optimal action a^* and pick as the share $\alpha^* = \sqrt{g(a^*)/E_{q^*(a^*)}[y]}$.

Moral Hazard with Asymmetric Ambiguity

Modify the symmetric model to incorporate the summary version of the following

Principal's information about technology is characterized as a single-valued mapping from a compact set of actions $\mathcal{A}$ to $\mathbb{K}_{\Delta(\mathcal{Y})}$ given by $Q^P(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$.

For each action $a \in \mathcal{A}$, Principal believes that this action induces a convex and compact set of distributions over outcomes denoted by $Q^P(a) \in \mathbb{K}_{\Delta(\mathcal{Y})}$.

Agent's information is characterized by $Q^A(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$.

Agent has more precise knowledge about technology than Principal: $Q^A(a) \subset Q^P(a)$ for each $a \in \mathcal{A}$. Moreover, Principal does not know the exact $Q^A(a)$.

Optimality of Linear Contracts under Asymmetric Ambiguity

Theorem (2)

Suppose $Q^A(a) \subset Q^P(a)$ *for all* $a \in \mathcal{A}$. *For any contract* $w(y)$, *there exists a linear contract that does as well, for both Principal and Agent.*

Summary and Further Questions

In summary, when robustness is sought in relation to pessimistic expectations for uncertain technology, as in the worst-case, linear contracts that align the contracting parties pessimistic expectations are optimal

- Allowing for risk aversion
- Under what conditions only linear contracts are optimal
- What about less extreme forms of pessimism?